Figure 1

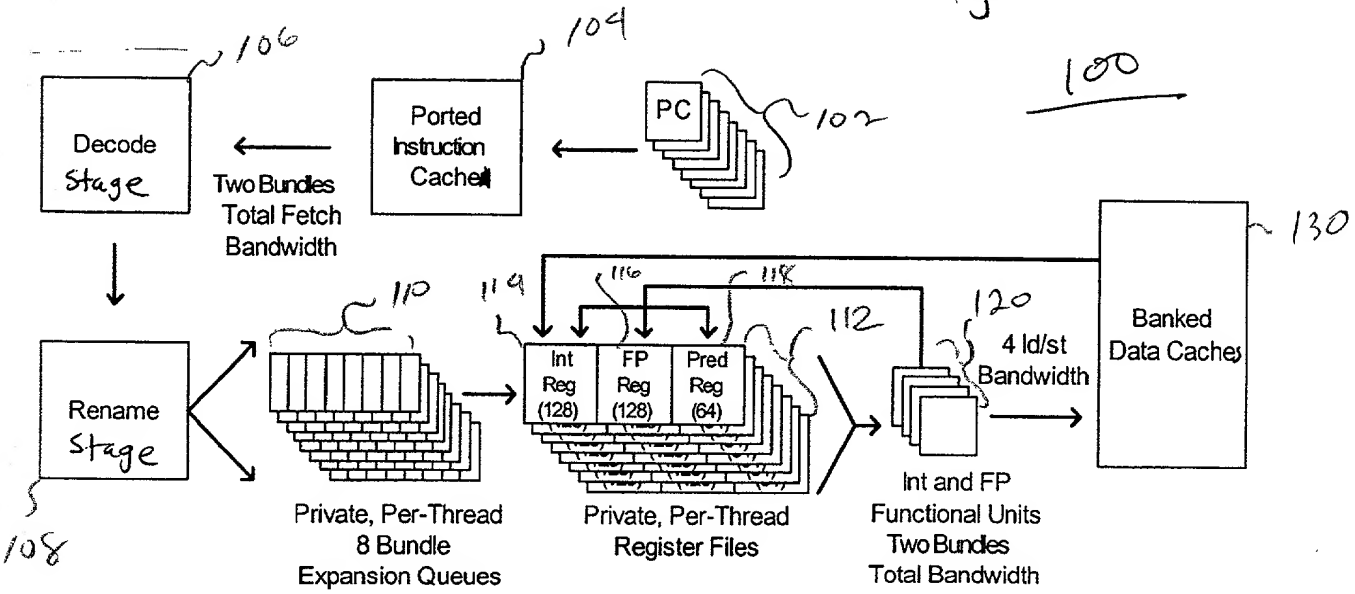
100
106
104
Decode Stage
Two Bundles Total Fetch Bandwidth
Ported Instruction Cache
PC
102
130
Banked Data Cache
110
114
116
118
112
120
4 ld/st Bandwidth
Rename Stage
108
Int Reg (128)
FP Reg (128)
Pred Reg (64)
Private, Per-Thread 8 Bundle Expansion Queues
Private, Per-Thread Register Files
Int and FP Functional Units Two Bundles Total Bandwidth

Figure 2

200

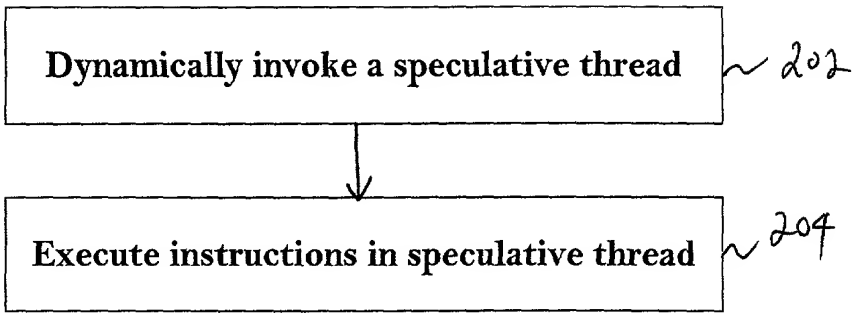


Figure 3

300

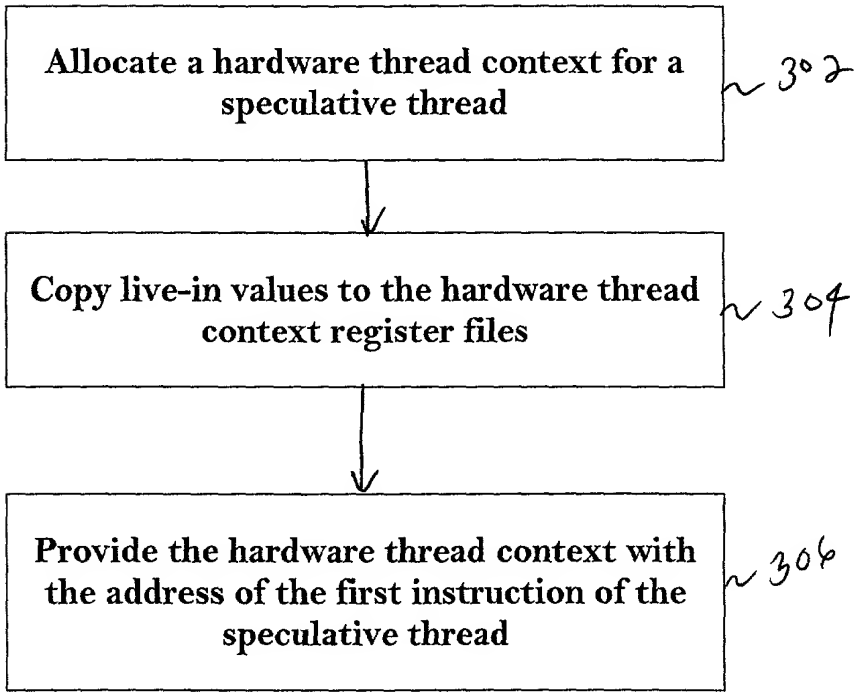

## Figure 4

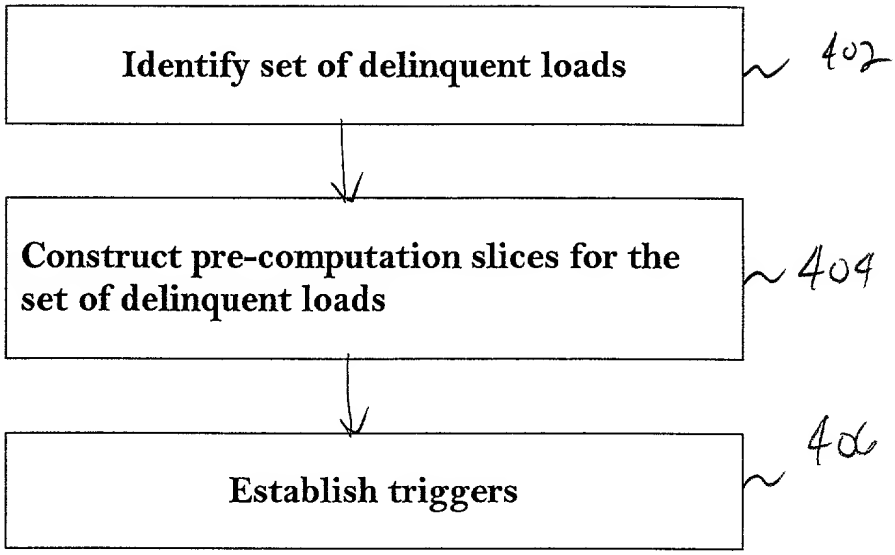


500

Figure 5

```
arc = arcs + group_pos;
for( ; arc < stop_arcs; arc += nr_group )
{
    if( arc ->ident  > BASIC )
    {
        red_cost = arc->cost - arc->    tail->potential   +
            arc-> head->potential  ;

        if((red_cost < 0 && arc->ident == AT_LOWER) ||
           (red_cost > 0 && arc->ident == AT_UPPER))
        {
            basket_size++;
            perm[basket_size]->a = arc;
            perm[basket_size]->cost = red_cost;
            perm[basket_size]->abs_cost = ABS(red_cost);
        }
    }
}
```

Delinquent Load 502

Delinquent Load 504

Delinquent Load 506

| | L1 Miss Rate / % Capacity Miss | L2 Miss Rate / % Capacity Miss | L3 Miss Rate / % Capacity Miss |
|---|---|---|---|
| DelinquentLoad 502 | 99.95% / 99.98% | 48.06% / 82.78% | 67.64% / 97.38% |
| DelinquentLoad 504 | 80.92% / 97.60% | 63.55% / 86.51% | 20.04% / 47.88% |
| DelinquentLoad 506 | 93.10% / 99.1% | 45.33% / 74.65% | 20.70% / 44.74% |

510

Figure 6

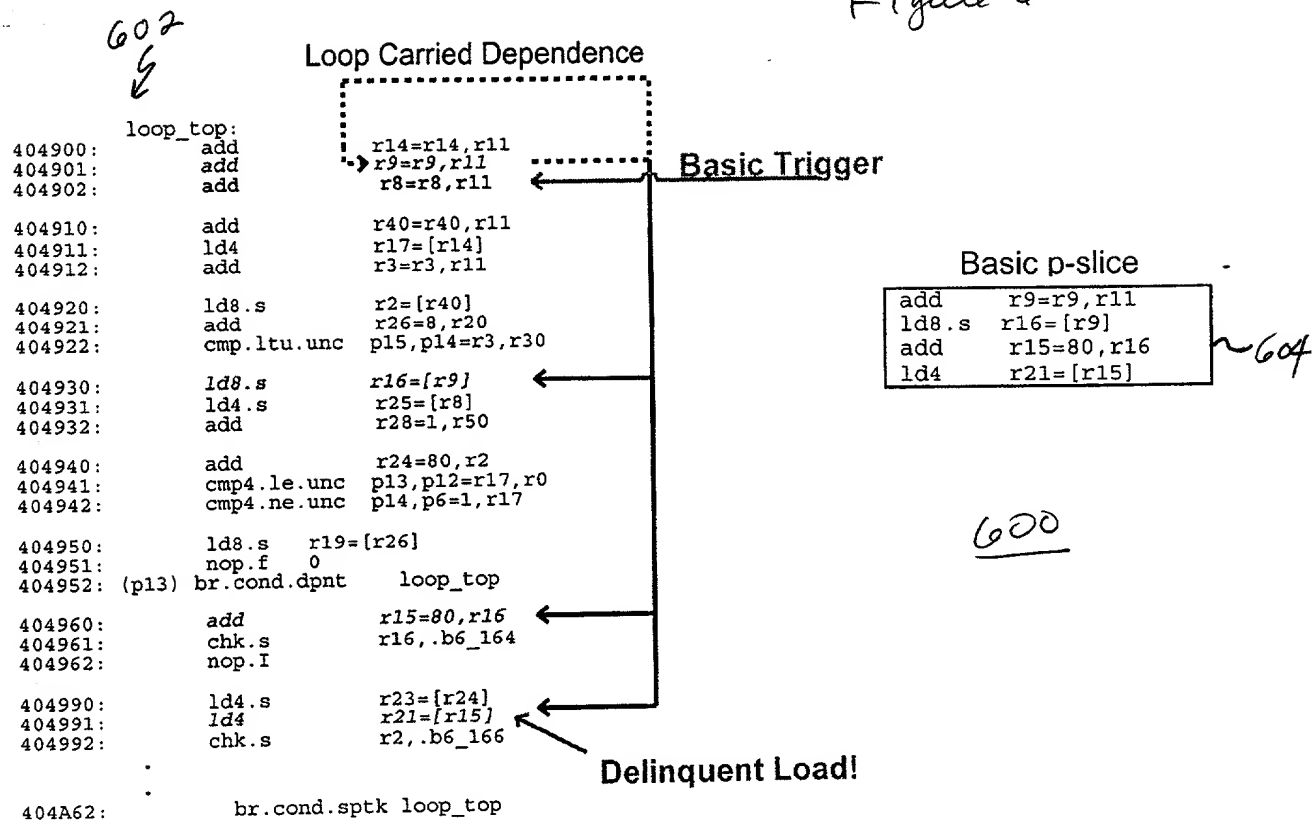


Figure 7

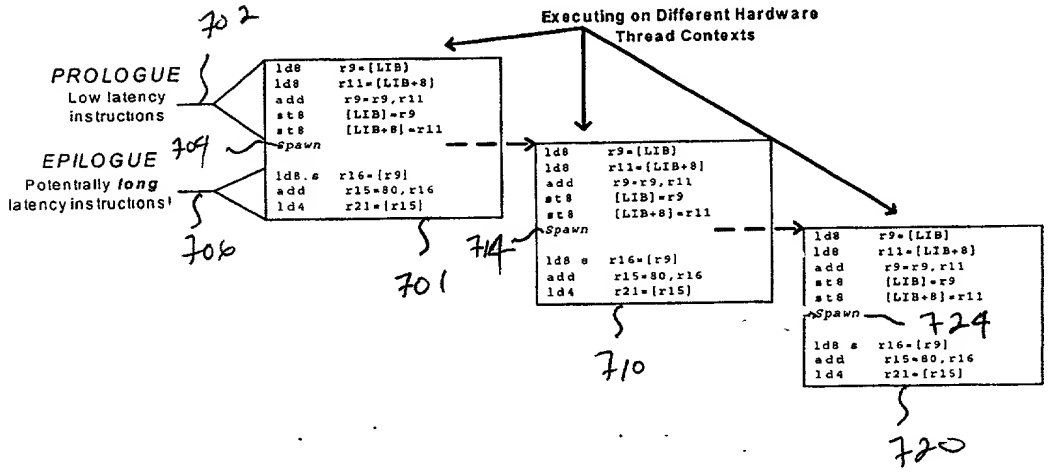

Figure 8

800

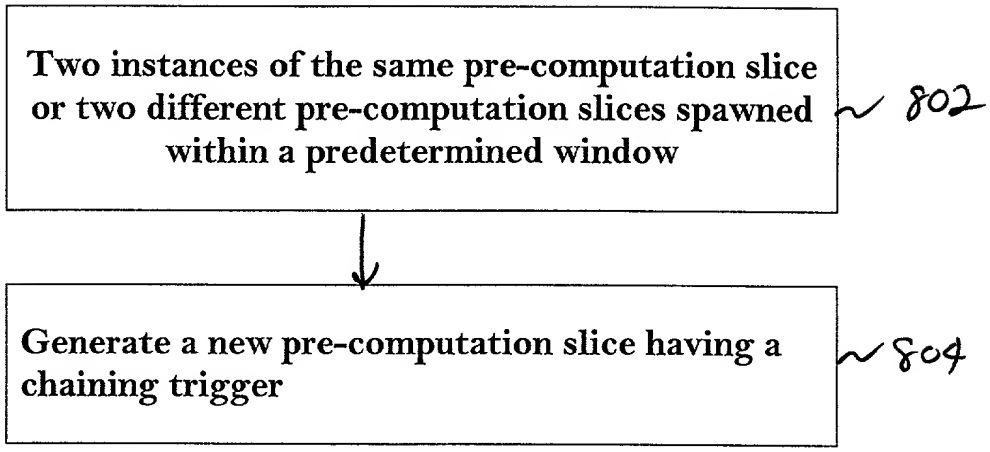


Figure 9

900

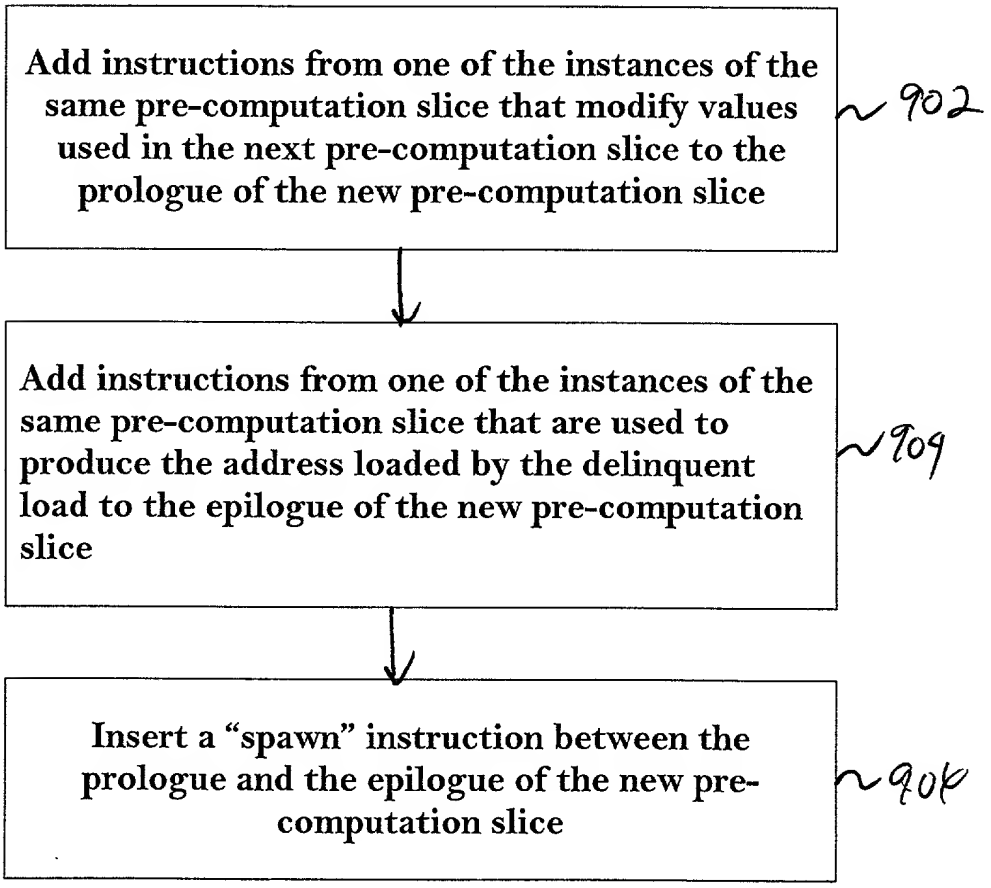

Figure 10

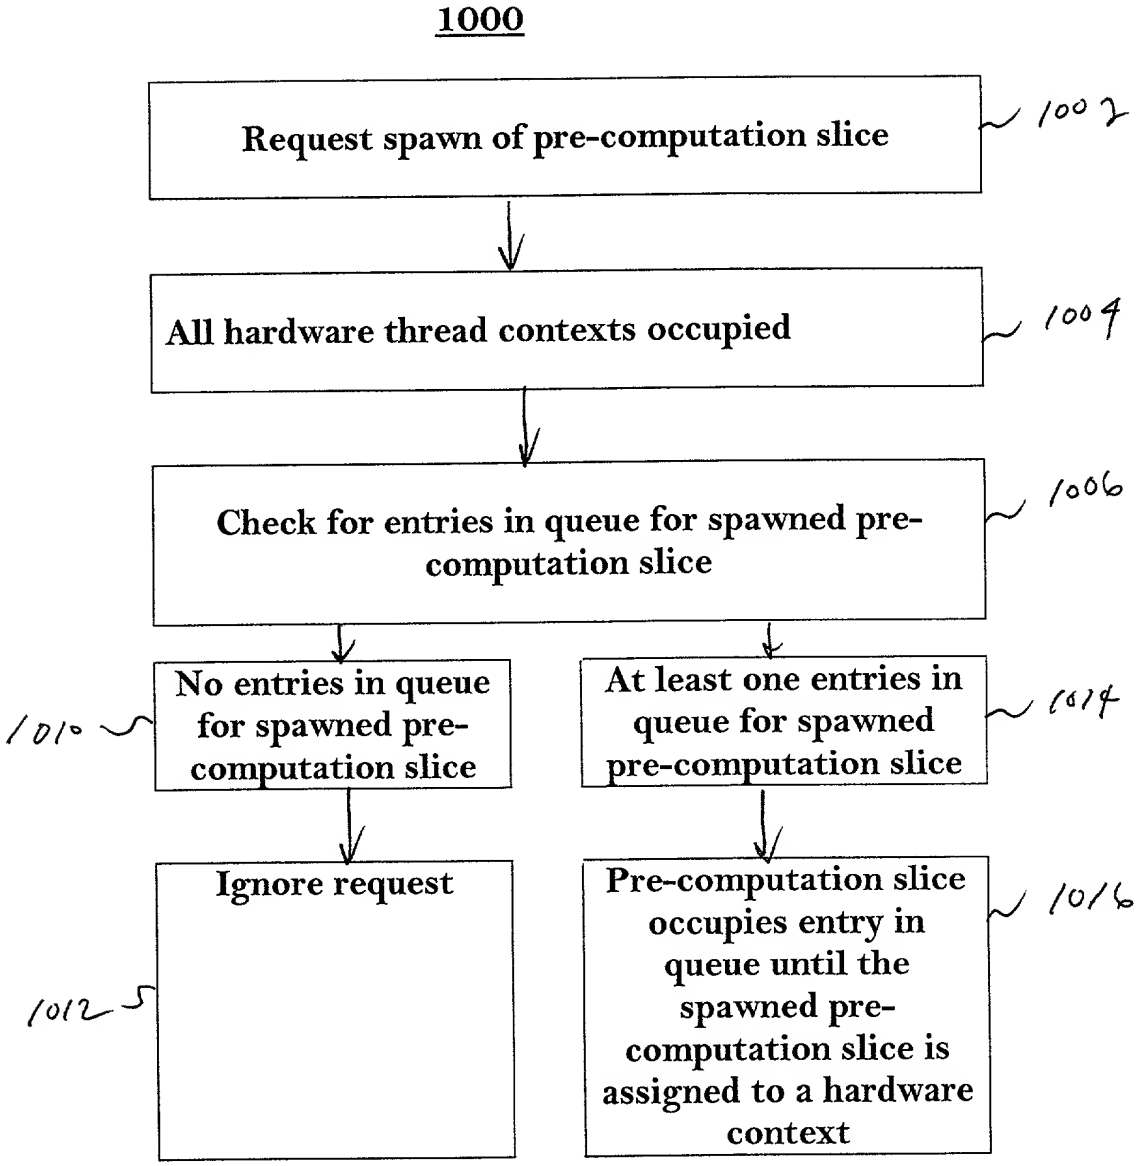
1000
Request spawn of pre-computation slice
1002
All hardware thread contexts occupied
1004
Check for entries in queue for spawned pre-computation slice
1006
No entries in queue for spawned pre-computation slice
1010
At least one entries in queue for spawned pre-computation slice
1014
Ignore request
1012
Pre-computation slice occupies entry in queue until the spawned pre-computation slice is assigned to a hardware context
1016